Federated Learning-Based Anomaly Detection for High-Performance Research Networks

Ehsan Saeedizade, Shamik Sengupta, Jay Thom
University of Nevada, Reno

Introduction

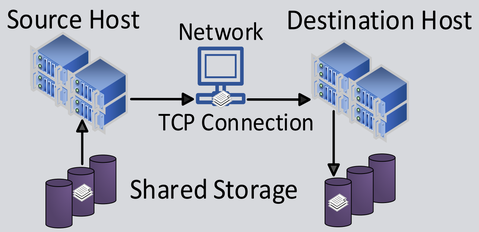

- High-Performance Research Networks (HPRNs) **enable large-scale scientific collaboration**
  - Cosmology SKA: **1 exabyte/13 days**
  - Genome sequencing: **2–40 exabytes/year**
- Challenges
  - Critical and **petabyte-scale transfers**
  - **Significant performance fluctuations** from interference
  - **Anomaly:** unexpected deviations in transfer performance
- Why it matters
  - Anomalies hurt reliability, throughput, performance

Prior Works and Gaps

- **(1) Centralized ML-based**
  - Raises **privacy risks**
  - Imposes **scalability challenges**
  - Fails to **generalize** across heterogeneous networks
- **(2) Federated Learning-based**
  - Applied to **IoT and cyber-attack detection**
  - Consider **IID data distribution**
  - Mostly focused on **binary classification** tasks

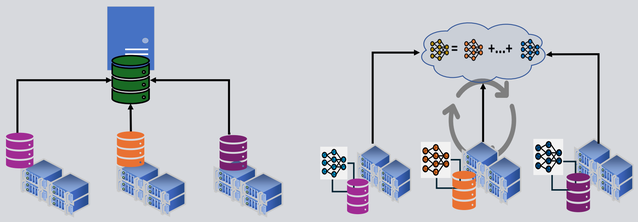

Proposed FL-based Approach

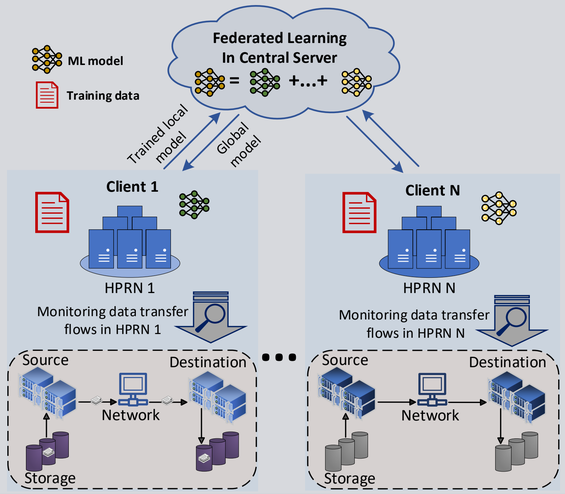

- **Cross-silo FL setup:** Each HPRN is treated as a federated client in a trusted collaboration.
- **Data Collection:** Each client monitors transfer performance metrics (144 total), covering sender, receiver, and network.
- **Feature Selection:** Random Forest with Mean Decrease in Impurity (MDI) identifies the 13 most critical features.
- **Transfer Learning:** Features are normalized relative to baseline values under normal conditions, creating domain-invariant representations that improve generalization across heterogeneous networks.
- **Federated Learning Process:** Local models are trained at each client, then aggregated on the central server using FedAvg. The global model is redistributed, repeating until convergence.

Key Results

- **Accuracy and F1-Score**
  FL + Transfer Learning (MLP) outperforms baselines
- **Convergence curve**
  Stabilizes in ~50 rounds (vs ~100 without TL)
- **Robustness in data-sparse environments**
  Clients detect missing anomaly in the trainset with >94% accuracy

Approach	MLP (Accuracy / F1)	XGBoost (Accuracy / F1)
FL	**96.6% / 0.966**	79.8% / 0.791
Federated-WTR	83.0% / 0.834	73.9% / 0.746
Centralized	79.2% / 0.799	**99.4% / 0.994**

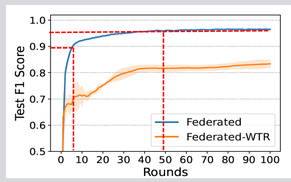

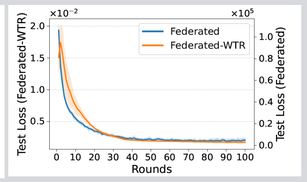

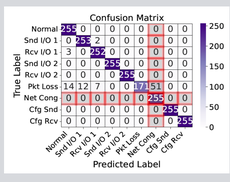

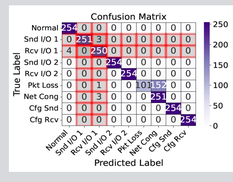

Contributions

🌐 **Framework** → FL-based anomaly detection for HPRNs
✨ **Transfer Learning** → improves generalization, faster convergence
🗄️ **Evaluation** → 8 heterogeneous testbeds, MLP & XGBoost
📊 **Performance** → 96.6% accuracy, F1 = 0.966, robust in sparse-data cases
🧠 **Model Insight** → MLP is better than XGBoost in FL

Future Work

🔀 **Multi-anomaly detection** → simultaneous anomalies
🖥️ **Client-side features** → local importance analysis
📏 **Severity-aware models** → incorporate anomaly severity

Acknowledgment | Contact

The work in this study was supported by the NSF under grant number #2346755.

References

[1] W. Marfo, D. K. Tosh, and S. V. Moore, “Network anomaly detection using federated learning,” in MILCOM 2022-2022 IEEE Military Communications Conference (MILCOM). IEEE, 2022, pp. 484–489.
[2] J. Zhao, S. Shetty, and J. W. Pan, “Feature-based transfer learning for network security,” in MILCOM 2017-2017 IEEE Military Communications Conference (MILCOM). IEEE, 2017, pp. 17–22.
[3] E. Saeedizade, B. Zhang, and E. Arslan, “Demystifying the performance of data transfers in high-performance research networks,” in 2023 IEEE 19th International Conference on e-Science (e-Science), 2023, pp. 1–11.